Social Discrete Choice Models

Danqing Zhang*
Civil & Environmental Engineering, UC Berkeley
danqing0703@berkeley.edu

Kimon Fountoulakis†
Statistics, UC Berkeley
kfount@berkeley.edu

Junyu Cao
Industrial Engineering & Operations Research, UC Berkeley
jycao@berkeley.edu

Mogeng Yin
Civil & Environmental Engineering, UC Berkeley
mogengyin@berkeley.edu

Michael Mahoney
Statistics, UC Berkeley
mmahoney@stat.berkeley.edu

Alexei Pozdnoukhov
Civil & Environmental Engineering, UC Berkeley
alexeip@berkeley.edu

arXiv:1703.07520v1 [cs.SI] 22 Mar 2017

KEYWORDS

Discrete choice models, computational social science, social influence, latent class model, Monte Carlo expectation maximization, alternating direction method of multiplier

1 ABSTRACT

Human decision making underlies data generating process in multiple application areas, and models explaining and predicting choices made by individuals are in high demand. Discrete choice models are widely studied in economics and computational social sciences. As digital social networking facilitates information flow and spread of influence between individuals, new advances in modeling are needed to incorporate social information into these models in addition to characteristic features affecting individual choices.

In this paper, we propose the first latent class discrete choice model that incorporates social relationships among individuals represented by a given graph. We add social regularization to represent similarity between friends, and we introduce latent classes to account for possible preference discrepancies between different social groups. Training of the social discrete choice models is performed using a specialized Monte Carlo expectation maximization algorithm. Scalability to large graphs is achieved by parallelizing computation in both the expectation and the maximization steps. Merging features and graphs together in a single discrete choice model does not guarantee success in terms of prediction. The data generating process needs to satisfy certain properties. We present and explain extensive experiments in which we demonstrate on which types of datasets our model is expected to outperform state-of-the-art models.

2 INTRODUCTION

In many application domains human decision making is modeled by discrete choice models. These models specify the probability that a person chooses a particular alternative from a given choice set, with the probability expressed as a function of observed and unobserved latent variables that relate to the attributes of the alternatives and the characteristics of the person. Multinomial logit models are in the mainstream of discrete choice models, with maximum likelihood used for parameter estimation from manually collected empirical data. It is important for practitioners to interpret the observed choice behaviors, and models that are linear in parameters are most common. At the same time, choice preferences within different social groups (though seemingly similar in terms of the observed characteristics) can vary significantly due to the unobserved factors or different context of the choice process. One way of accounting for this is to introduce latent class models. Latent class logistic regression models are common tools in multiple domains of social science research [4, 15, 18].

It is also recognized that social influence can be a strong factor behind variability in choice behaviors. The impact of social influence on individual decision-making has attracted a lot of attention. Researchers have employed laboratory experiments, surveys, and studied historical datasets to evaluate the impact of social influence on individual decision making. It is however difficult to avoid an identification problem in the analysis of influence processes in social networks [12]. One has to account for endogeneity in explanatory variables in order for claims of causality made by these experiments to be useful [8]. Due to these limitations of observational studies of influence, randomized controlled trials are becoming more common. In general, distinguishing social influence in decision making from homophily, which is defined as the tendency for individuals with similar characteristics and choice behaviors to form clusters in social networks, is currently a growing area of research and debate [3, 16].

Data science research in the area is more concerned with developing models that scale to large graphs. Social connectivity information is used to improve predictive performance of the models, with auto-regressive approaches based on label propagation being widespread. Graph regularization ideas that penalize parameter differences among the connected nodes were studied in the context of classification, clustering and recommender systems [1, 11], with recent advances in distributed optimization applied to parametric models on networks [9].

In this paper, we introduce social graph regularization ideas into latent class discrete choice modeling. We aim to combine the expressiveness of parametric model specifications and descriptive exploratory power of latent class models with recent advances in distributed optimization for parameter estimation. Our model specification is introduced in Section 3. Parameter estimation of the social discrete choice models is performed using a specialized Monte Carlo expectation maximization algorithm presented in Section 4. We adopt modern optimization techniques to achieve scalability to large graphs by parallelizing computation in both the expectation

*equally contributing author
†equally contributing author

and the maximization steps. Finally, in Section 5, we present and explain extensive experiments in which we demonstrate on which types of social processes and datasets our model is expected to provide advatages over the state-of-the-art methods. We also describe our approach of dealing with missing labels at nodes that can not be removed without altering the social structure of the graph. We then draw some conclusions and outline future work, including extensions to deep models, in Section 6.

3 SOCIAL MODELS

In this section we discuss how one can add social factors in choice models and the advantages of doing so. We start by explaining notation. Then we will provide certain examples and comment on cases where data are missing and how one can treat this problem.

We define $[N] := \{1, 2, \ldots, N\}$, $i \in [N]$, $t \in [K]$, where N and K are integers. We will use the following notation and definitions.

- N: number of individuals,
- K: number of latent classes,
- n: number of samples per node,
- d: number of features for each individual,
- $x_i \in \mathbb{R}^{d \times n}$: feature-samples matrix of individual i,
- $z_i \in [K]$: latent class variable of individual i,
- $y_i \in \{-1, 1\}$: binary choice of individual i ,
- $W_t \in \mathbb{R}^d$: model coefficients of class t,
- $b_{it} \in \mathbb{R}$: model offset coefficients of individual i with class t,
- $\mathcal{V}$: set of nodes in a social graph, with each node corresponding to an individual,
- $\mathcal{E}$[1]: set of edges, presenting relationship between two individuals, $(i, j) \in \mathcal{E}$ means that there exists an edge between individuals i and j.

We assume that there is only one sample per node, i.e., $n = 1$, however, the proposed models can be extended to the case of $n > 1$. We further assume that the graph $(\mathcal{V}, \mathcal{E})$ is unweighted, noting that the models can be extended to weighted graphs. In the following subsections we occasionally drop indices i and t depending on the context to simplify notation. We denote with $\theta := \{W, b\}$ the set of model coefficients W_t, b_{it}, $\forall i, t$.

Let $h_{it}(x_i) := W_t^T x_i + b_{it}$, we consider the following choice model for individual i of class t,

$$y_{it} = \begin{cases} 1 & \textbf{if } h_{it}(x_i) \geq c, \\ -1 & \textbf{otherwise,} \end{cases} \tag{1}$$

where $c \in \mathbb{R}$ is a decision threshold constant.

3.1 Social logistic regression

Choice model specified by Eq. (1) includes several known models such as logistic regression, where $K = 1$ and y_i follows a Bernoulli distribution given x_i. To incorporate the social aspect in logistic regression one assumes that the parameters b follow an exponential family parametrized with the given graph

$$P(b) \propto \prod_{(i,j)\in\mathcal{E}} e^{-\lambda(b_i - b_j)^2},$$

[1]We assume that each edge in the graph is represented in the set $\mathcal{E}$ only once.

where $\lambda \in \mathbb{R}$ is a hyper-parameter. This model is usually trained by using a maximum log-likelihood estimator which reduces to the following regularized logistic regression problem

$$\theta^* := \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^{N} \log\left(1 + e^{-y_i h_i(x_i)}\right) + \lambda \sum_{(i,j)\in\mathcal{E}} (b_i - b_j)^2,$$

where $h_i(x_i) := W^T x_i + b_i$. Notice that the social information, i.e., edges $\mathcal{E}$, appears as Laplacian regularization for the coefficients b.

3.2 Latent class social logistic regression

Our first extension is a logistic regression with latent classes, $K > 1$. In this model, y_{it} follows a Bernoulli distribution given x_i and $z_i = t$. To incorporate social information we assume that latent class variables z_i are distributed based on the following exponential family parametrized by the given social graph

$$P(z; b) \propto \prod_{(i,j)\in\mathcal{E}} \exp\left(-\lambda \sum_{t=1}^{K} (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t)\right),$$

where b represents the collection of coefficients b_{it} $\forall i, t$ which are the parameters of the distribution. In this model we assume that each individual has its one local coeffficient b_{it} for each class. Notice that this model does not penalize different coefficients b_{it} among connected individuals in different classes. This is because we assume that connected individuals in different classes should have independent linear classifiers.

In the common latent class models, latent class variables are independent and identically distributed following the multinomial distribution, i.e., $z_i \sim \text{Mult}(\pi, 1)$ $\forall i$, where π is the probability of success. However, in our specification hidden variables are correlated and are not identically distributed. Note that this specification allows studying social structures that underlie an observed choice process, and hence extends the range of inferences possible with the state-of-the-art models.

A graphical interpretation of this model is given in Figure 1. The resulting model can be trained using maximum likelihood and the Expectation-Maximization (EM) algorithm; details are discussed in Section 4.

4 PARAMETER ESTIMATION

We will use the maximum likelihood estimator to estimate the parameters W and b for the social choice models. We will discuss how to train the parameters of the models for the case where there is more than one class, i.e., $K > 1$. To train models with latent classes it is common to use the Expectation-Maximization (EM) algorithm. In this section we provide details for the Expectation step (E-step) and the Maximization step (M-step).

4.1 Monte Carlo EM

Correlation among latent variables imposed by the social graph do not allow exact calculation of posterior distributions in the E-step using standard EM approaches. Instead, an approximate calculation of the E-step using Monte Carlo EM (MCEM) [5, 10, 13, 20] is employed. It is a modification of the original EM algorithm where the E-step is conducted approximately using a Monte Carlo Markov Chain (MCMC) algorithm. The details for each step of the MCEM

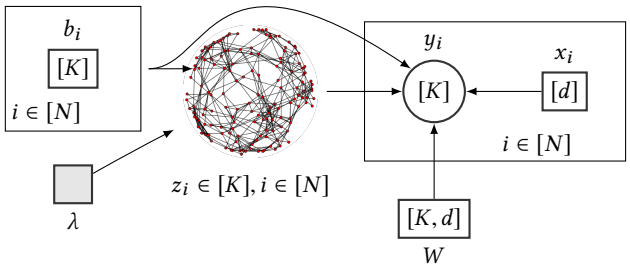

Figure 1: Graphical model representation for social logistic regression models with latent variables. We use a modified plate notation to represent conditional dependence among random variables and dependence on parameters. In particular, random variables are represented using circles and their number is shown in brackets inside the circle, i.e., y_i corresponds to K variables. Parameters are represented in rectangles and their size is shown in brackets with two components, i.e., W corresponds to $K \times d$ coefficients. Data are shown in rectangles and their size in brackets, i.e., x_i corresponds to d features. There are N nodes in the graph and each node corresponds to a random variable z_i which takes values in $[K] := \{1, \ldots, K\}$. The hyper-parameter λ is represented using a grey rectangle.

algorithm for the proposed models are provided in the following subsections.

Algorithm 1 MCEM algorithm for social choice models

1: **Inputs:**
   $(x_i, y_i), i = 1 \ldots N$
2: **Initialize:**
   $\theta^0 := \{W^0, b^0\} \leftarrow$ arbitrary value, $k \leftarrow 0$
3: **repeat**
4: **E-step:** (Subsection 4.1.1)
5: Calculate approximate node posterior
6: for each node $i \in [N]$

$$q(z_i = t) := P(z_i = t | y_i, x_i; b^k)$$

, and for each edge $(i, j) \in \mathcal{E}$, the edge posterior

$$q(z_i = z_j = t) := P(z_i = t, z_j = t | y_i, y_j, x_i, x_j; b^k)$$

7: by using the block MCMC sampling (Algorithm 2).
8: **M-step:** (Subsection 4.1.2)
9: Solve the optimization problem

$$\theta^{k+1} := \operatorname*{argmin}_{\theta} Q(\theta; x, y),$$

   where $Q(\theta; x, y)$ is defined at (4).
10: $k \leftarrow k + 1$
11: **until** termination criteria are satisfied.

4.1.1 Expectation step. In this step the objective is to compute the marginal posterior distribution for nodes and edges, which will be used in the M-step to calculate the negative expected log-likelihood function. Please refer to the Appendix for derivation of negative expected log-likelihood, which reveals the need for calculation of marginal posterior distributions.

In particular, for the E-step one needs to calculate the following node marginal posterior probability

$$P(z_i = t | y_i, x_i; \theta) = \frac{P(y_i | x_i, z_i = t; \theta) P(z_i = t; b)}{\sum_{s=1}^{K} P(y_i | x_i, z_i = s; \theta) P(z_i = s; b)} \tag{2}$$

and the following edge posterior probability

$$P(z_i = t, z_j = t | y_i, y_j, x_i, x_j; \theta) = \frac{P(y_i, y_j | x_i, x_j, z_i = t, z_j = t; \theta) P(z_i = t, z_j = t; b)}{\sum_{m,q=1}^{K} P(y_i, y_j | x_i, x_j, z_i = m, z_j = q; \theta) P(z_i = m, z_j = q; b)}, \tag{3}$$

where θ represents the collection of parameters W and b. For small graphs we can approximate the above distributions using standard MCMC algorithms. For large graphs please refer to the Scalability of the E-step Subsection 4.2.1.

4.1.2 Maximization step. Let us denote with $q(z_i = t) = P(z_i := t | y_i, x_i; \theta)$ and $q(z_i = z_j = t) := P(z_i = t, z_j = t | y_i, y_j, x_i, x_j; \theta)$ the marginal posterior distributions. The M-step of the EM algorithm requires minimizing the negative expected log-likelihood function

$$Q(\theta; x, y) := \sum_{i \in V} \sum_{t=1}^{K} q(z_i = t) \log \left(1 + e^{-y_i h_{it}(x_i)}\right) + \lambda \sum_{(i,j) \in E} \sum_{t=1}^{K} (b_{it} - b_{jt})^2 q(z_i = z_j = t), \tag{4}$$

where θ represents the collection of parameters W and b and $h_{it}(x_i) := W_t^T x_i + b_{it}$. Derivation of this function is given in Subsection 7.1 in the Appendix. For small graphs standard convex optimization solvers can be used. For large graphs please refer to Section 4.2.2 where we discuss how we can maximize the expected log-likelihood efficiently with a distributed algorithm.

4.2 Scalability

In this subsection we discuss how to distribute computation to make our algorithm flexible enough to deal with large graphs.

4.2.1 Scalability of E-step. Calculating the marginal posterior probabilities (2) and (3) is computationally expensive due to the marginal probabilities $P(z_i = t; b)$ and $P(z_i = t, z_j = t; b)$. This is because to calculate the latter two we have to marginalize $N - 1$ and $N - 2$ latent variables, respectively. We avoid this problem by using a block MCMC sampling technique to compute $P(z_i = t; b)$ and $P(z_i = t, z_j = t; b)$. The algorithm is given in Algorithm 2.

The algorithm uses a preprocessing step to partition the graph into c disjoint communities. Then it runs an MCMC algorithm on each community/block in *parallel* by ignoring the edges among the blocks.

Let us comment on how the number of blocks c affects convergence of the MCMC Algorithm 2 to $P(z_i = t; b)$ and $P(z_i = t, z_j = t; b)$ in practice. In Figure 2 we present the relative error of the MCMC algorithm for various values of c, settings of parameters

Algorithm 2 Block MCMC sampling

1: **Inputs:**
   Graph with nodes $\mathcal{V}$ and edges $\mathcal{E}$, parameters θ, number of blocks c
2: **Preprocess:** Use community detection algorithm, i.e., [14], to partition the graph in c *disjoint* communities/blocks, i.e., $C := \{C_1, \ldots, C_c\}$, $\bigcup_{q=1}^{c} C_q = \mathcal{V}$ and $\bigcap_{q=1}^{c} C_q = \emptyset$.
3: **parfor** each community in C **do**
4: Compute the node and edge marginal probabilities
5: $P(z_i = t; b)$ and $P(z_i = t, z_j = t; b)$, respectively, using
6: an MCMC sampling algorithm by taking into account
7: edges only among nodes in the community.
8: **end parfor**

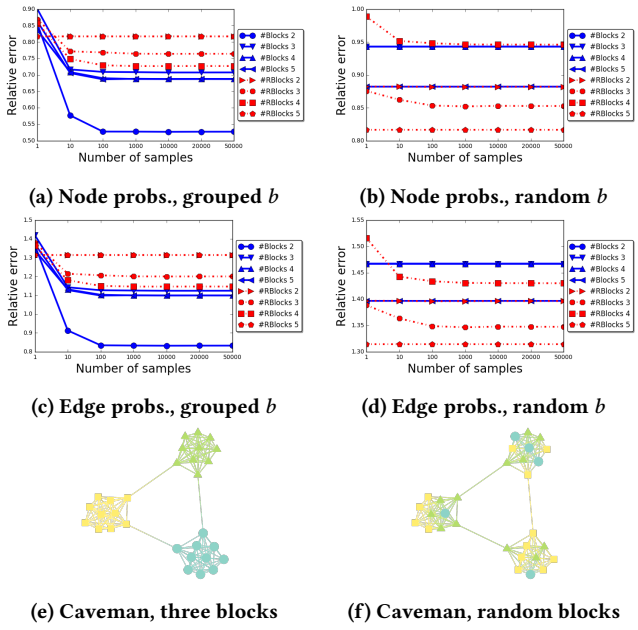

(a) Node probs., grouped b **(b) Node probs., random b**

(c) Edge probs., grouped b **(d) Edge probs., random b**

(e) Caveman, three blocks **(f) Caveman, random blocks**

Figure 2: This figure illustrates the accuracy of blocked MCMC Algorithm 2. Details are provided in Section 4.2.1.

b and ways of grouping nodes for the caveman graph [19]. The caveman graph consists of three communities of ten nodes each, see Figure 2e. We consider two classes for the latent variables. We make use of two ways of grouping nodes, the first uses a community detection technique [14], see Figure 2e, the second is by grouping nodes randomly, see Figure 2f. Moreover, we consider two settings of b. 1) Parameters b are separated in three groups that correspond to the three groups of nodes of the graph. For each group parameters b follow a Gaussian distribution with different mean such that the parameters are significantly different among groups but similar within the groups. Figures 2a and 2c correspond to the first setting of parameters b. 2) Parameters b follow the same Gaussian distribution regardless of the groups of the graph. Figures 2b and 2d correspond to the second setting of parameters b. The two settings correspond to the last and the first iterations of the MCEM algorithm, respectively, since we expect that close to stationary points of the likelihood function parameters b are grouped based on the communities of the given graph while at the first iterations parameters b have a random setting; we discuss this in Section 5.

Figures 2a and 2c illustrate accuracy of block MCMC algorithm for node and edge probabilities, respectively. Notice that the accuracy of block MCMC is overall better when community detection is used compared to random assignment of nodes in communities. This is because preprocessing helps to minimize the edges among the blocks that are ignored. On the other hand, in Figures 2b and 2d when parameters b are not grouped, then community detection does not perform better than random assignment of nodes in communities. Moreover, notice in Figures 2b and 2d that the performance of block MCMC does not improve as the number of samples increases. This happens because parameters b are randomly generated and they are not grouped, then based on MCMC updating rule, we get that $P(z_i = t; b) \approx 1$ for $t = 0$ and $P(z_i = 1; b) \approx 0$ for $t = 0$.

Let us now comment briefly on the theoretical asymptotic convergence of MCEM to a stationary point of the likelihood function. Convergence theory of MCEM in [5, 13] states that if Algorithm 2 is set to work with one block $c = 1$ (no partition of the graph), and the MCMC sample size increases deterministically across MCEM iterations, then MCEM converges almost surely. A consequence of blocking of latent variables for the MCMC algorithm is that asymptotic convergence of MCEM is not guaranteed anymore. However, in practice often MCEM is terminated without even knowing if the algorithm converges to an accurate solution. See for example Section 5 in [13] and references therein about arbitrary termination criteria of MCEM. Therefore, we consider that the parallelism of block MCMC Algorithm 2 offers a trade-off among convergence and computational complexity per iteration by controlling the number of blocks, which in practice can speed up each iteration of the MCEM algorithm significantly.

4.2.2 Scalability of M-step. In this section we discuss how to minimize in distributed manner the negative expected log-likelihood function of the M-step of MCEM Algorithm 1. Following the work of [9] that applies ADMM to network lasso method, we extend it by allowing both local and global variables on nodes in the presence of latent classes. We will focus the derivation on the example of social logistic regression with latent classes which was discussed in Subsection 3.2. For reasons of simplifying the notation we do not provide details for the social neural network model, but we do not comment at appropriate places how ADMM can be applied for this model.

First notice that the negative expected log-likelihood function (4) can be separated in K independent problems, one for each class t. Therefore, in this section we discuss how to apply distributed ADMM for each problem. Let

$$Q(\theta; x, y, t) := \sum_{i \in V} q(z_i = t) \log\left(1 + e^{-y_i h_{it}(x_i)}\right) + \lambda \sum_{(i,j) \in E} (b_{it} - b_{jt})^2 q(z_i = z_j = t) \tag{5}$$

be the objective function for problem with latent class t, where θ represents the collection of parameters W and b and $h_{it}(x_i) := W_t^T x_i + b_{it}$. Then $Q(\theta; x, y)$ in (4) is equal to $\sum_{t=1}^{K} Q(\theta; x, y, t)$. Let us define $q_{it} := q(z_i = t)$ and $q_{ijt} := q(z_i = z_j = t)$. To minimize

(5) for a given class t using ADMM we introduce a copy of b_{it} denoted by z_{ij}, $\forall i$ and a copy of W_t denoted by g_i, $\forall i$,

$$\begin{aligned} \min_{W,b} \quad & \sum_{i\in V} q_{it} \log\left(1 + e^{-y_i h_{it}(x_i)}\right) + \lambda \sum_{(i,j)\in\mathcal{E}} q_{ijt}(z_{ij} - z_{ji})^2 \\ \text{s.t:.} \quad & b_{it} = z_{ij}\ j \in \mathcal{N}(i),\ \forall i \\ & W_t = g_i,\ \forall i, \end{aligned} \tag{6}$$

where $\mathcal{N}(i)$ are the adjacent nodes of node i. By introducing copies for b_{it} $\forall i$ we dismantle the sum over edges into separable functions, additionally, by introducing copies for W_t, we dismantle the sum over the nodes for the logistic function. Then by relaxing the constraints we can make the problem (6) separable which opens the door to distributed computation.

We define the augmented Lagrangian below, where u and h are the dual variable and ρ_1 and ρ_2 are the penalty parameters.

$$\begin{aligned} L_\rho(W_t, b_t, g, z, u, h; t) := \sum_{i\in V} \Big\{ & q_{it} \log\left(1 + e^{-y_i(g_i^T x_i + b_{it})}\right) \\ & + \frac{\rho_1}{2}\left(\|h_i\|_2^2 + \|W_t - g_i + h_i\|_2^2 \right) \Big\} + \lambda \sum_{(i,j)\in\mathcal{E}} \Big\{ q_{ijt}(z_{ij} - z_{ji})^2 \\ & + \frac{\rho_2}{2}\left(\left\|u_{ij}\right\|_2^2 + \left\|u_{ji}\right\|_2^2 + \left\|b_{it} - z_{ij} + u_{ij}\right\|_2^2 + \left\|b_{jt} - z_{ji} + u_{ji}\right\|_2^2 \right) \Big\}. \end{aligned}$$

The resulting ADMM algorithm is presented in Algorithm 3, where $f(z_{ij}, z_{ji}) := L_\rho(W_t^{k+1}, b_t^{k+1}, g^{k+1}, (z_{ij}, z_{ji}, z_{(ij)^c}^k), u^k, h^k; t)$. Notice that the subproblems in Step 6 do not have closed form solutions, however, they can be solved efficiently using an iterative algorithm since they are univariate problems which depend only on x_i and not all data. Similarly, the subproblems in step 9 do not have closed form solution, but they have only d unknown variables and depend only on x_i and not all data. Moreover, Step 12 has a closed form solution, which corresponds to solving a 2×2 linear system. Observe that the ADMM algorithm 3 can be run in a distributed setting by distributing the data among processors, as the subproblems in Steps 6 and 9 depend only on x_i and not all data.

Algorithm 3 ADMM for problem 6

1: **Initialize:**
   $k \leftarrow 0$, W^k, b^k, g^k, z^k, u^k and h^k
2: **repeat**
3: Set $W_t^{k+1} = \frac{1}{N}\sum_{i=1}^N (g_i^k - h_i^k)$
4: $b_{it}^{k+1} := \arg\min_{b_{it}} L_\rho(W_t^{k+1}, b_{it}, g^k, z^k, u^k, h^k; t)\ \forall i \in \mathcal{V}$
5: $g_i^{k+1} := \arg\min_{g_i} L_\rho(W_t^{k+1}, b_t^{k+1}, g_i, z^k, u^k, h^k; t)\ \forall i \in \mathcal{V}$
6: $z_{ij}^{k+1}, z_{ji}^{k+1} = \arg\min_{z_{ij}, z_{ji}} f(z_{ij}, z_{ji})\ \forall (i,j) \in \mathcal{E}$
7: Set
$$h_i^{k+1} = h_i^k + (W_t^{k+1} - g_i^{k+1})\ \forall i \in \mathcal{V}$$
$$u_{ij}^{k+1} = u_{ij}^k + (b_{it}^{k+1} - z_{ij}^{k+1})\ \forall (i,j) \in \mathcal{E}$$
$$u_{ji}^{k+1} = u_{ji}^k + (b_{jt}^{k+1} - z_{ji}^{k+1})$$
8: $k \leftarrow k + 1$
9: **until** termination criteria are satisfied.

5 EXPERIMENTS

In this section, we analyze the empirical performance of the proposed social model on a range of datasets. We outline practical recommendations and illustrate examples where the proposed model is most suitable.

Our implementation reproducing the below results is available at https://github.com/DanqingZ/social-DCM. The number of iterations of Gibbs sampler in the E-step grows with the number of iterations of the MCEM algorithm. The M-step is implemented using ECOS solver [7] embedded in CVXPY [6] for W and b updates within ADMM iterations.

5.1 Synthetic data

We demonstrate how having choices separable on the given graph helps to improve prediction when social information are taken into account in the following two ways. i) Varying the connectivity between different classes in the graph. ii) Varying the discrepancies of choices in communities of the graph. In what follows, when we mention separability of choices on the graph we mean that the decisions y_i are clustered based on communities in the graph. By separable in the feature space we mean that there exist linear hyperplanes which separate the data x_i accurately.

5.1.1 Varying connectivity between classes. We consider $N = 300$ individuals, each node corresponds to one feature vector x_i, $d = 10$ and $y_i \in \{-1, 1\}$ $\forall i$. 150 feature vectors $x_i \in \mathbb{R}^d$ are generated using a Gaussian distribution where each component in x_i follows $G(\mu = 10, \sigma = 0.1)$, and the rest 150 feature vectors have components which follow $G(\mu = 30, \sigma = 0.1)$. We randomly split the individuals into three communities with 100 individuals each. We assume that there are two classes, i.e., $K = 2$, and that the individuals in two out of three communities are in the first class, i.e., $t = 1$, while the individuals in the remaining community are in the second class, i.e., $t = 2$. The communities are randomly assigned to classes. All individuals in the first class have $y_i = 1$, while all individuals in the second class have $y_i = -1$. Notice that the x_i vectors are assigned to communities regardless of their Gaussian distribution and y_i is set based on the classes, which correspond to communities. Therefore, the choices y_i are separable on the graph but not separable in the feature space. We set the probability of two individuals that are in the same community to get connected to 0.2. We set the probability of two individuals that are in the same class but not in the same community to get connected to 0.01. We set the probability of two individuals in different classes to get connected to β, which is the parameter that controls the connectivity between classes.

Figure 3 shows the graph structure when $\beta = 10^{-4}$, $\beta = 10^{-2}$ and $\beta = 10^{-1}$ Notice that the larger β is the more edges among the communities that belong in different classes. Table 1 shows the prediction result of four models as a function β. Notice that since x_i and y_i are not changed as β changes, therefore, the prediction of logistic regression and logistic regression with latent class remains constant at 62%. The reason that these models perform poorly is because the choices y_i are not separable given the feature vectors x_i only. Observe in Table 1 that when β is as small as 10^{-4}, which means that individuals in different classes are very unlikely to get connected, see Figure 3a, then the prediction result of the proposed

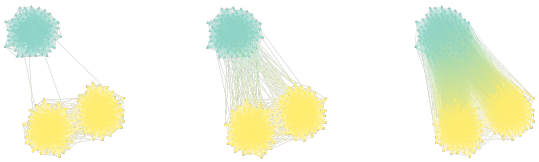

(a) $\beta = 10^{-4}$ (b) $\beta = 10^{-2}$ (c) $\beta = 10^{-1}$

Figure 3: The nodes with square shape and yellow color correspond to class $t = 1$. The nodes with triangle shape and turquoise color correspond to class $t = 2$. The larger β the more edges among nodes with different class.

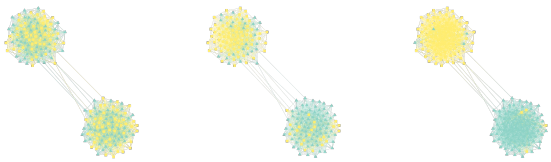

(a) $\|w\|_2 = 0$ (b) $\|w\|_2 = 4$ (c) $\|w\|_2 = 6$

Figure 4: Three synthetic examples showing the influence of parameter w on class preference. The nodes with square shape and yellow color correspond to choice $y_i = 1$. The nodes with triangle shape and turquoise color correspond to choice $y_i = -1$.

social models is larger than 80%. On the other hand, when β becomes larger, the prediction of the social models is declining. But as long as $\beta < 0.1$, the proposed social model with latent class, see Subsection 3.2, is performing better than logistic regression and logistic regression with latent class models. When $\beta = 0.1$ the social models has the same prediction performance as the logistic regression and latent class models. This is because the classes are not clearly separable on the graph, see Figure 3c.

Table 1: Prediction results on a randomly chosen test set of 50 individuals when β is varied, i.e., connectivity between classes. For all models the regularization parameter λ which corresponds to the best prediction result out of a range of parameters is chosen.

	β				
model	10^{-4}	10^{-3}	5×10^{-3}	10^{-2}	10^{-1}
logistic reg.	62%	62%	62%	62%	62%
log. reg. lat. class	62%	62%	62%	62%	62%
social no lat. class	80%	62%	62%	62%	62%
social with lat. class	88%	82%	64%	62%	62%

5.1.2 Varying choice preference parameters. An ideal scenario for the proposed social models is when classes correspond to communities of the given graph and when the choices y_i are clustered according to the classes. However, choices y_i might be misplaced in wrong classes. We study how the preference difference between classes affects the performance of the proposed model.

For this experiment individual feature vectors x_i are generated from a Gaussian distribution for a set of $N = 200$ individuals, each component of x_i follows $G(\mu = 1, \sigma = 0.1)$. We randomly split this set of individuals into two equal parts, each part represents a class where individuals share the same parameters w. Let us assume the w in the first class is w_1 and the w in the second class is $w_2 = -w_1$. In our experiments we vary parameter w_1. For each individual b_i satisfies $b_i \sim G(0, 0.1)$. For the graph setting, we set the probability of people in the same class to be connected, and the probability of people in different classes to be connected to 0.2 and 10^{-4}, respectively. This way, we can make sure classes correspond to communities.

Based on data generation process, by tuning $\|w_1\|$, we are able to get full control of preference difference among individuals in the two classes. When $\|w_1\|$ becomes larger then preference difference becomes larger as well. As we can see in Figure 4 when $\|w_1\|$ becomes larger more individuals in class one have $y_i = 1$ (i.e., yellow squares) and more individuals in class two have have $y_i = 1$ (i.e., turquoise triangles). When $w = 0$ in both classes around half of the individuals have $y_i = 1$, the other half $y_i = -1$, which means there is no preference difference between the two classes. Prediction results for this experiment are shown in Table 2.

Table 2: Prediction results on a randomly chosen test set of 50 individuals when $\|w\|_2$ is varied. For all models the regularization parameter λ which corresponds to the best prediction result out of a range of parameters is chosen.

	$\|w\|_2$					
model	10	5	3	2	1	0
logistic reg.	48%	44%	42%	52%	58%	52%
log. reg. lat. class	48%	48%	54%	54%	42%	36%
social with lat. class	100%	100%	94%	86%	68%	36%

5.2 Adolescent smoking

This example uses a dataset collected by [2]. This research program, known as the teenage friends and lifestyle study, has conducted a longitudinal survey of friendships and the emergence of the smoking habit (among other deviant behaviours) in teenage students across multiple schools in Glasgow, Scotland.

5.2.1 Dataset. Social graphs of 160 students (shown in Figure 5) within the same age range of 13-15 years is constructed following a surveyed evidence of reciprocal friendship, with an edge placed among individuals i and j if individual i and individual j named each other friends. We included five variables into the feature vector x_i: age; gender; money: indicating how much pocket money the student had per month, in British pounds; romantic: indicating whether the student is in a romantic relationship; family smoking: indicating whether there were family members smoking at home. Notice that the feature vectors x_i and the edges of the graph are different at different timestamps. The response variable y represents the stated choice is whether a student smokes tobacco ($y_i = 1$), otherwise $y_i = -1$.

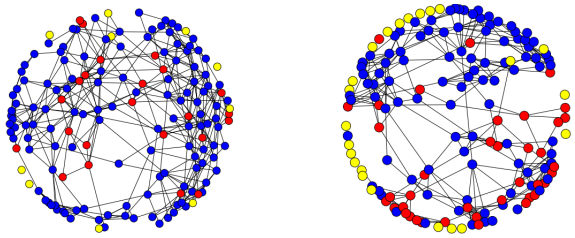

(a) Start of the study (Feb 1995): 127 non-smokers (blue), 23 smokers (red), 10 unobserved (yellow) 422 edges.

(b) End of the study (Jan 1997): 98 non smokers (blue), 39 smokers (red), 23 unobserved (yellow) 339 edges.

Figure 5: Social graphs of student friendships and smoking behaviors within the 2 years period of the study.

5.2.2 Missing labels. In practice one might have full observation of the graph structure, but not being able to observe choices y_i for each individual. Let us denote $\tilde{\mathcal{V}}$ as the set of nodes for which one observes choices y_i, and $\bar{\mathcal{V}}$ as the set of nodes for which the choices are not revealed. One approach to model specification in this case would be to remove nodes $\bar{\mathcal{V}}$ and their edges from the graph, disregarding data x_i. However, here we empirically illustrate that it is preferable to keep nodes $\bar{\mathcal{V}}$, their edges and their coefficients b_i in the model and simply set their unrevealed choice to $y_i = 0$, $\forall i \in \bar{\mathcal{V}}$. The impact of this specification of the model on parameter estimation can be illustrated as follows. For example, for social logistic regression with $K = 1$ the maximum log-likelihood problem reduces to

$$\theta^* := \underset{\theta}{\operatorname{argmin}} \sum_{i \in \tilde{\mathcal{V}}} \log\left(1 + e^{-y_i h_i(x_i)}\right) + \lambda \sum_{(i,j) \in \mathcal{E}} (b_i - b_j)^2,$$

where in the first sum one ignores nodes in $\bar{\mathcal{V}}$, but maintains in the second sum their edges and coefficients b_i.

Maintaining nodes with b_i $\forall i \in \bar{\mathcal{V}}$, and their edges in the graph correspond to additional regularization among nodes, since two nodes that are connected through a node with no observed choice are forced indirectly to have similar coefficients b_i. It is particularly important in applications where there are strong reasons to believe that social influence is a significant factor driving the choice of the individuals. To provide further intuitions, consider a thought experiment of comparing the structure of the social graph before and after removing nodes in $\bar{\mathcal{V}}$ (shown yellow in Figure 5) and their edges. Our experiments suggest that maintaining the density of the graph helps avoiding overfitting and results in estimating models with more accurate out-of-sample predictions.

5.2.3 Models comparison. We measured predictions on this dataset and report here 3-fold cross validation results of four models: i) logistic regression; ii) latent class logistic regression; iii) social logistics regression, see Subsection 3.1; iv) social latent class logistics regression, see Subsection 3.2. Cross-validation process treats the removed nodes as nodes with missing labels, as described above. We consider $K = 2$ latent classes in this experiment. The prediction performance is shown in Tables 3-4.

Table 3: Adolescent smoking prediction, February 1995

model	accuracy
logistics regression	81.1%
latent class logistics regression	78.9%
social logistics regression	80.0%
social latent class logistics regression	82.2%

Table 4: Adolescent smoking prediction, January 1997

model	accuracy
logistics regression	68.5%
latent class logistics regression	72.1%
social logistics regression	65.5%
social latent class logistics regression	77.1%

High performance of logistic regression and latent class model for the beginning of the study (Table 3) indicates that the smoking preferences are largely defined by the objective factors. Moreover, parameters in the underlying classes of the latent class model don't differ much. Social latent class model performs equally well. This difference grows significantly larger when a confounding variable of smoking in the family is removed from the feature list.

Furthermore, by the end of the study (January 1997, Table 4) one can see a significantly higher predictive accuracy of the social latent class logistics regression model. It may indicate that smoking within a certain group of teenagers have become a social norm (indicating the difference in offset parameters b_{it}), or that the response y_i to independent factors x_i within that group differs from the others as reflected by differences in W_t.

Notice the significant decrease of prediction accuracy for non social models among Tables 3 and 4. This is because at the beginning of the study (February 1995) less than 15% of teenagers smoked, while at the end of the study (January 1997) about 25% of teenagers smoked. This difference is likely due to social norms developed among teenagers which are captured by the graph and therefore missed by non social models.

5.2.4 Parameter visualization. We are going to illustrate how the specification of the proposed model allows in-depth exploration of the parameters to assist in making this type of conclusions. To that end, we are going to explore parameters b_{it} and class membership probabilities across the regularization path of hyper-parameter λ.

The estimated class membership (probability of being in a given latent class) on the graph is shown in Figure 6. A visualization of estimated b_{it} for one latent class for several values of λ is shown in Figure 7. When $\lambda = 0.1$, smoking pattern begins to show across the graph, i.e., compare Figures 5b and 7b. Let us note that $\lambda = 0.1$ value corresponds to the best prediction accuracy in our experiments for the proposed social latent class logistic regression model for the smoke data at the end of the study (Jannuarry 1997). This is because the social latent class logistic regression model is able to clearly distinguish a group of socially connected individuals within which the choice preferences towards smoking are higher.

When $\lambda = 0.01$, b_{it} are similar across the nodes, one the other hand, when $\lambda = 10$, b_{it} are not similar across the nodes. Although

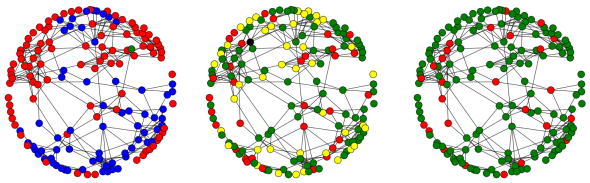

(a) $\lambda = 10$ **(b)** $\lambda = 0.1$ **(c)** $\lambda = 0.01$

Figure 6: Class membership probabilities estimated for the nodes for multiple values of λ. Blue, yellow, green, black and red colours correspond to probabilities about 0, $(0, 0.5)$, 0.5, $(0.5, 1)$ and 1, respectively.

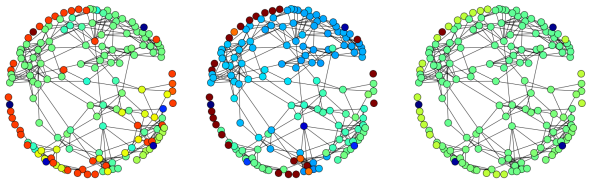

(a) $\lambda = 10$ **(b)** $\lambda = 0.1$ **(c)** $\lambda = 0.01$

Figure 7: The values of b_{it} estimated for the nodes for multiple values of λ. Lighter color corresponds to higher values.

this is counter-intuitive since the graph regularization favours similar b_{it} across nodes for large values of λ, it is explained by the node and edge posterior distribution in the M-step which also controls regularization across nodes during the execution of the algorithm.

6 CONCLUSIONS AND FUTURE WORK

In this paper we introduced social graph regularization ideas into latent class discrete choice modeling. We have developed, implemented, and explored parameter estimation algorithms that allow parallel processing implementation for both E- and M-steps, benefiting from recent advances in distributed optimization based on ADMM methods. In experimental evaluation, we have focused on investigating the usefulness of the model in revealing and supporting hypothesis in studies where not only predictive performance (that was found to be highly competitive), but also understanding social influence is crucial. Our model can be directly applied to study social influence on revealed choices in large social graphs with rich node attributes (that are very rarely available in open access due to privacy issues).

Another promising direction for future work is as follows. With the recent work of [17], the specification we propose can accommodate decision functions of multiple layers L conditional on the latent class $z_i \in [K]$, $L > 1$ and $K \geq 1$. Similarly to the "shallow" models specification, to incorporate social information we assume that the variables z_i are distributed based on an exponential family parametrized by the given graph, but for this extension we take into account multiple layers, introducing non-linearity in choice

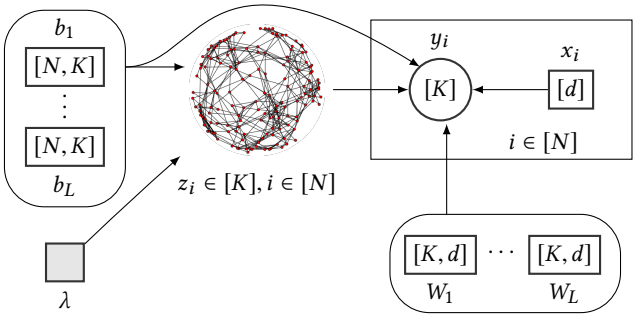

Figure 8: Graphical model representation for a deep social choice model with latent variables. In addition to notation in Figure 1, it includes rounded rectangles which denote groups of coefficients.

dependency on x_i. In particular, we assume that

$$P(z; b) \propto \prod_{(i,j)\in\mathcal{E}} \exp\left(-\lambda \sum_{l,t=1}^{L,K} (b_{itl} - b_{jtl})^2 \mathbf{1}(z_i = z_j = t)\right),$$

where b represents the collection of coefficients b_{itl} $\forall i, t, l$. The graphical representation of this model is shown in Figure 8. The major differences and an advantage of the deep social choice model over the social latent class logistic regression model introduced in Subsection 3.2 is the non-linearity of the decision y_{it} in (1). As shown in Figure 8, the structure of the graphical model remains the same, except for that the additional layers result in an increasing the number of b and W coefficients that needs to be estimated using a maximum likelihood and the Monte Carlo expectation-maximization algorithm. In the M-step, one has to introduce copies of W and b for each layer in order to make the objective function separable [17]).

7 APPENDIX

7.1 Negative expected log-likelihood in Eq. (4)

We denote with $q(z) := P(z|y, x; \theta)$ the posterior distribution, with $\sum_z$ the sum over all latent variables z, and θ represents the collection of parameters W and b. Let $\mathbf{1}(z_i = t)$ be the indicator function, which is equal to 1 if $z_i = t$. We assume that latent variables follow the following distribution:

$$P(z; b) \propto \prod_{(i,j)\in\mathcal{E}} \exp\left(-\lambda \sum_{t=1}^{K} (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t)\right),$$

and

$$P(y_i|x_i, z_i = t, \theta) = \frac{1}{1 + e^{-y_i h_i(x_i)}}, \tag{7}$$

where $h_{it}(x_i) := W_t^T x_i + b_{it}$. The derivation of the expected log-likelihood is shown below.

$$\begin{aligned}
\tilde{Q}(\theta; x, y) &:= \sum_z q(z) \log P(y, z|x; \theta) \\
&= \sum_z q(z) \log(P(y|x, z; \theta) P(z; b)) \\
&= \sum_z q(z) \log P(y|x, z; \theta) + \sum_z q(z) \log P(z; b) \\
&= \sum_z q(z) \log \left(\prod_{i \in \mathcal{V}} \sum_{t=1}^{K} P(y_i|x_i, z_i = t, \theta) \mathbf{1}(z_i = t) \right) \\
&\quad - \lambda \sum_z q(z) \sum_{(i,j) \in E} \sum_{t=1}^{K} (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t) \\
&= \sum_z q(z) \sum_{i \in \mathcal{V}} \log \left(\sum_{t=1}^{K} P(y_i|x_i, z_i = t, \theta) \mathbf{1}(z_i = t) \right) \\
&\quad - \lambda \sum_z q(z) \sum_{(i,j) \in E} \sum_{t=1}^{K} (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t).
\end{aligned}$$

We can exchange the sequence of log and $\sum_{t=1}^{K}$ because each node can only be in one class, thus we have

$$\begin{aligned}
\tilde{Q}(\theta; x, y) &= \sum_z q(z) \sum_{i \in \mathcal{V}} \sum_{t=1}^{K} \mathbf{1}(z_i = t) \log P(y_i|x_i, z_i = t; \theta) \\
&\quad - \sum_z q(z) \sum_{(i,j) \in E} \sum_{t=1}^{K} \lambda (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t).
\end{aligned}$$

Let's write the summation over z inside the summation over vertices and the summation over latent variables

$$\begin{aligned}
\tilde{Q}(\theta; x, y) &= \sum_{i \in \mathcal{V}} \sum_{t=1}^{K} \sum_z q(z) \mathbf{1}(z_i = t) \log P(y_i|x_i, z_i = t; \theta) \\
&\quad - \lambda \sum_{(i,j) \in E} \sum_{t=1}^{K} \sum_z q(z) (b_{it} - b_{jt})^2 \mathbf{1}(z_i = z_j = t),
\end{aligned}$$

and then we get the marginal probabilities

$$\begin{aligned}
\tilde{Q}(\theta; x, y) &= \sum_{i \in \mathcal{V}} \sum_{t=1}^{K} \log(P(y_i|x_i, z_i = t, \theta)) \sum_z q(z) \mathbf{1}(z_i = t) \\
&\quad - \lambda \sum_{(i,j) \in E} \sum_{t=1}^{K} (b_{it} - b_{jt})^2 \sum_z q(z) \mathbf{1}(z_i = z_j = t) \\
&= \sum_{i \in \mathcal{V}} \sum_{t=1}^{K} q(z_i = t) \log P(y_i|x_i, z_i = t, \theta) \\
&\quad - \lambda \sum_{(i,j) \in E} \sum_{t=1}^{K} (b_{it} - b_{jt})^2 q(z_i = z_j = t).
\end{aligned}$$

Using (7) and multiplying by minus equation $\tilde{Q}(\theta; x, y)$ we get the negative expected log-likelihood function in (4).

REFERENCES

[1] J. Abernethy, O. Chapelle, and C. Castillo. 2010. Graph regularization methods for web spam detection. *Mach. Learn.* 81, 2 (2010), 207–225.

[2] H Bush, Patrick West, and Lynn Michell. 1997. The role of friendship groups in the uptake and maintenance of smoking amongst pre-adolescent and adolescent children: Distribution of Frequencies. *Glasgow, Scotland: Medical Research Council* (1997).

[3] Nicholas A Christakis and James H Fowler. 2013. Social contagion theory : examining dynamic social networks and human behavior. November 2011 (2013). DOI : http://dx.doi.org/10.1002/sim.5408

[4] H. Chung, B. P. Flaherty, and J. L. Schafer. 2006. Latent class logistic regression: application to marijuana use and attitudes among high school seniors. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 169, 4 (2006), 723–743.

[5] B. Deylon, M. Lavielle, and E. Moulines. 1999. Convergence of a Stochastic Approximation Version of the EM Algorithm. *The Annals of Statistics* 27, 1 (1999), 94–128.

[6] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.

[7] Alexander Domahidi, Eric Chu, and Stephen Boyd. 2013. ECOS: An SOCP solver for embedded systems. In *Control Conference (ECC), 2013 European*. IEEE, 3071–3076.

[8] Elenna Dugundji and Joan Walker. 2005. Discrete Choice with Social and Spatial Network Interdependencies: An Empirical Example Using Mixed Generalized Extreme Value Models with Field and Panel Effects. *Transp. Res. Rec.* 1921, 1 (2005), 70–78. DOI : http://dx.doi.org/10.3141/1921-09

[9] David Hallac, Jure Leskovec, and Stephen Boyd. 2015. Network lasso: Clustering and optimization in large graphs. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 387–396.

[10] R. A. Levine and G. Casella. 2012. Implementations of the Monte Carlo EM Algorithm. *Journal of Computational and Graphical Statistics* 10, 3 (2012), 422–439.

[11] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. 2011. Recommender systems with social regularization. *WSDM '11 Proceedings of the fourth ACM international conference on Web search and data mining* (2011), 287–296.

[12] Charles F. Manski. 1993. Identification of Social Endogenous Effects: The Reflection Problem. *Rev. Econ. Stud.* 60, 3 (1993), 531–542. DOI : http://dx.doi.org/10.2307/2298123

[13] Ronald C Neath and others. 2013. On convergence properties of the Monte Carlo EM algorithm. In *Advances in Modern Statistical Theory and Applications: A Festschrift in Honor of Morris L. Eaton*. Institute of Mathematical Statistics, 43–62.

[14] T. P. Peixoto. 2014. Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models. *Phys. Rev. E* 89 (2014), 012804.

[15] K. Roeder, K. G. Lynch, and D. S. Nagin. 1999. Modeling Uncertainty in Latent Class Membership: A Case Study in Criminology. *J. Amer. Statist. Assoc.* 94, 447 (1999), 766–776.

[16] Cr Shalizi and Ac Thomas. 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociol. Methods Res.* (2011), 1–27. arXiv:arXiv:1004.4704v3 http://smr.sagepub.com/content/40/2/211.short

[17] Gavin Taylor, Ryan Burmeister, Zheng Xu, Bharat Singh, Ankit Patel, and Tom Goldstein. 2016. Training neural networks without gradients: A scalable admm approach. In *International Conference on Machine Learning*.

[18] J. L. Walker and J. Li. 2007. Latent lifestyle preferences and household location decisions. *Journal of Geographical Systems* 9, 1 (2007), 77–101.

[19] D. J. Watts. 1999. Networks, Dynamics, and the Small-World Phenomenon. *Amer. J. Soc.* 105 (1999), 493–527.

[20] G. C. .G. Wei and M. A. Tanner. 1990. A Monte Carlo Implementation of the EM Algorithm and the Poor Man's Data Augmentation Algorithms. *J. Amer. Statist. Assoc.* 85, 411 (1990), 699–704.